

PIRATTE: Proxy-based Immediate Revocation of ATtribute-based Encryption

Sonia Jahid and Nikita Borisov
 {sjahid2,nikita}@illinois.edu
 University of Illinois at Urbana-Champaign

Abstract—Access control to data in traditional enterprises is typically enforced through reference monitors. However, as more and more enterprise data is outsourced, trusting third party storage servers is getting challenging. As a result, cryptography, specifically Attribute-based encryption (ABE) is getting popular for its expressiveness. The challenge of ABE is revocation. To address this challenge, we propose PIRATTE, an architecture that supports fine-grained access control policies and dynamic group membership. PIRATTE is built using attribute-based encryption; a key and novel feature of our architecture, however, is that it is possible to remove access from a user without issuing new keys to other users or re-encrypting existing ciphertexts. We achieve this by introducing a proxy that participates in the decryption process and enforces revocation constraints. The proxy is minimally trusted and cannot decrypt ciphertexts or provide access to previously revoked users. We describe the PIRATTE construction and provide a security analysis along with performance evaluation. We also describe an architecture for online social network that can use PIRATTE, and prototype application of PIRATTE on Facebook.

Index Terms—Attribute-based Encryption, Revocation, Access Control, Social Networking

1 INTRODUCTION

Access control to data in traditional enterprises is typically provided by reference monitors that enforce a particular policy. This approach, however, creates a vulnerability for monitors that are buggy or compromised and thus do not enforce the correct policy. This is a particular concern in large, distributed enterprises, as well as Online Social Networks (OSNs), where recent privacy compromises showcase the dangers of entrusting privacy controls to social network providers [1], [2]. The problem gets challenging as enterprises outsource more and more data, and rely on third party storage servers to enforce the access policy.

Attribute-based encryption (ABE) [3], [4], [5] provides an alternative approach to data protection, where the ability to decrypt data items is controlled by a policy specified in terms of attributes. ABE systems mimic the expressiveness of traditional access control systems, but use cryptography instead of reference monitors. In enterprises, this means that data can be sent over channels and stored on media that might at some point become compromised without fear of a privacy breach. Among several existing schemes, Ciphertext Policy ABE (CP-ABE) [4] is appropriate for most applications. In CP-ABE, the encryptor uses some public parameters and a policy described over attributes to encrypt a piece of data. Different secret keys are issued for different sets of attributes. A key that has enough attributes to satisfy the policy decrypts the ciphertext.

Various applications can benefit from the flexibility

and expressiveness of ABE, but require the support of frequent revocation. ABE falls short in such areas when frequent and immediate revocation of access is required. Researchers have proposed revocation by attaching an expiry date to the keys [4], [5] or introducing proxies [6]. However, existing approaches come with their shortcomings either by introducing delay in revocation, increasing the size of ciphertext, or affecting (re-keying) all the users including both the revoked and non-revoked ones.

In this paper, we present PIRATTE, a proxy-based immediate revocation scheme for CP-ABE. Our design makes use of a minimally trusted proxy, which handles revoked users and attributes. Upon revocation, no new key is generated for any user, neither is the existing data re-encrypted. We believe this feature is key for access control in any context where ABE is used together with highly dynamic group membership and large datasets. Note that the proxy is minimally trusted: it cannot decrypt by itself, and even if it were compromised, it cannot allow previously revoked users to decrypt either. The only assumption we hold is that the proxy is updated with a new key each time a revocation takes place. PIRATTE ensures forward secrecy, backward secrecy with some assumptions, immediate revocation of complete or partial access, and delegation of access with single and multiple key authorities.

Our Contribution:

- We provide the construction of our scheme including Proxy-based key/attribute revocation and

access delegation, and the security analysis of the scheme.

- We implement a prototype named PIRATTE and compare its performance with the CP-ABE scheme by Bethencourt et al. [4] (BSW CP-ABE).
- In addition, we describe a case study that can benefit from using PIRATTE. We choose OSN since recent research in OSNs proposes the use of cryptography to enhance privacy [7], [8], [9]. However, they are not completely successful because of some shortcomings of existing cryptographic schemes. We also present an application of PIRATTE running on Facebook platform.

Roadmap:

The rest of the paper is organized as follows. We briefly discuss background information on CP-ABE and the base revocation scheme in Section 2. Next, we provide a detailed description of our construction in Section 3. Section 4 discusses the security of our construction. We describe the applicability of PIRATTE to OSNs in Section 5. We describe our performance analysis and the Facebook application in Section 6, related work in Section 7, and conclude in Section 8.

2 BACKGROUND

In this section we describe some background information necessary to understand our scheme. We describe two cryptographic schemes that form the foundation of our approach.

2.1 Attribute-Based Encryption

Cryptography can be used to enforce access control to information by encrypting data such that only authorized users can decrypt it. However, with standard public-key cryptography, it is necessary to explicitly enumerate all of the users who may decrypt each data item. While it is possible to issue group keys, this creates complex key management issues and several problems still remain.

Attribute-based encryption provides a better solution for defining fine-grained access control to groups of people. With ABE, a user Alice assigns sets of attributes to other users (we will call these users Alice’s contacts) and issues the corresponding secret keys. She encrypts data items using policies expressed in terms of plain attributes that defines what attributes one must have in order to decrypt the piece of data (this variant is called *ciphertext-policy* ABE; *key-policy* ABE reverses the process, with attributes assigned to data items and policies to keys).

ABE can support complex policies, such as “(Friend OR Co-worker) AND Neighbor”. ABE also explicitly prevents collusion between users: if Bob is Alice’s co-worker, and Carol is her neighbor, they cannot combine their attributes together to satisfy the above policy if neither of them satisfies it individually. Finally,

ABE provides public-key functionality, allowing, for example, Bob to encrypt a message with Alice’s public key to be decrypted by her contacts to whom Alice assigns secret attribute keys.

We can formally define a CP-ABE scheme by four algorithms with an option for attribute delegation [4]:

- **SETUP**. This algorithm takes security parameters and generates a public key PK and master secret key MK .
- **ENCRYPT**(PK, M, P). This algorithm takes the public key PK , a message M , and a policy P and generates a ciphertext CT encrypted with P .
- **KEYGEN**(MK, S). This algorithm uses the master secret key MK to generate a secret attribute key SK using the attributes in the set S .
- **DECRYPT**(CT, SK). This algorithm decrypts a ciphertext CT to plaintext M as long as the set of attributes S in SK satisfies the policy P that was used to generate CT from M . (The policy is implicitly encoded in CT .)
- **DELEGATE**(SK, \tilde{S}). The delegate algorithm takes as input a secret key SK for some set of attributes S and a set $\tilde{S} \subseteq S$. It outputs a secret key \tilde{SK} for the set of attributes \tilde{S} .

2.2 Revocation Scheme

To support practical revocation in PIRATTE, we adapt the broadcast revocation scheme developed by Naor and Pinkas to prevent digital piracy [10]. The scheme uses Shamir secret sharing [11] to create shares of a secret key, where $t + 1$ shares are necessary for reconstruction, and gives one share to each user.

During regular operation, the distributor broadcasts t random shares, which lets any user to reconstruct the secret key by combining the shares with his or her own. To revoke up to t users, the distributor broadcasts their shares instead of the random ones. Any non-revoked user still has $t + 1$ distinct shares and can reconstruct the secret, whereas the revoked users do not have enough information even if they all collude.

3 CONSTRUCTION

3.1 Assumptions and Basics

Before going into the details of the construction, we present some basic mathematical assumptions, and the details of CP-ABE and the revocation scheme used in PIRATTE.

Bilinear Pairing

Let \mathbb{G}_1 , \mathbb{G}_2 , and \mathbb{G}_T be multiplicative cyclic groups of prime order p , and e a map ($\mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$). Let g_1 and g_2 be generators of \mathbb{G}_1 and \mathbb{G}_2 respectively ($\mathbb{G}_i = \langle g_i \rangle$). If $\forall u \in \mathbb{G}_1, v \in \mathbb{G}_2$ and $a, b \in \mathbb{Z}_p$, $e(u^a, v^b) = e(u, v)^{ab}$

and $e(g_1, g_2) \neq 1$, then e is called a bilinear pairing. If $\mathbb{G}_1 = \mathbb{G}_2$, it is called a symmetric pairing, otherwise the pairing is asymmetric.

Secret Sharing

In Shamir's secret sharing scheme [11], a secret s in some field F is shared among n parties by creating a random polynomial $P \in F[x]$ of degree t such that $P(0) = s$. The i -th party gets the share $\langle i, P(i) \rangle$. Given any $t + 1$ shares $P(x_0), \dots, P(x_t)$, it is possible to recover $P(0)$ using Lagrange interpolation:

$$P(0) = \sum_{i=0}^t \lambda_i P(x_i), \quad \text{where } \lambda_i = \prod_{j \neq i} \frac{x_j}{(x_j - x_i)}$$

BSW CP-ABE Scheme

The algorithms in CP-ABE due to the Bethencourt et al. are described below. Though CP-ABE uses symmetric pairing, it can be implemented using an asymmetric pairing as well.

- **SETUP:** The key authority KA generates a public key PK, and a master secret key MK:

$$PK = \mathbb{G}_1, g, h = g^\beta, e(g, g)^\alpha$$

$$MK = (\beta, g^\alpha)$$

$$\text{where random } \alpha, \beta \in \mathbb{Z}_p, \mathbb{G}_1 = \langle g \rangle, |\mathbb{G}_1| = p$$

The PK also contains an extra component $f = g^{1/\beta}$ to support attribute delegation.

- **ENCRYPT(PK, M, τ):** A policy is represented as an access tree structure τ with the attributes at leaves and threshold k -of- n gates at intermediate nodes. Each node is associated with a polynomial q_x of degree d_x , where d_x is 1 less than the threshold value k of that node. The polynomials are of degree 0 for OR gates and leaves. The secret s (random $s \in \mathbb{Z}_p$) to blind the data M is associated with the polynomial at the root of the tree, i.e., $q_R(0) = s$. The sharing works in a top down manner: for all other nodes, $q_x(0) = q_{\text{parent}(x)}(\text{index}(x))$. $\text{index}(x)$ returns a number between 1 and num associated with x where num is the number of children of $\text{parent}(x)$. Let Y be the set of leaf nodes in τ . The ciphertext CT is:

$$CT = (\tau, \tilde{C} = Me(g, g)^{\alpha s}, C = h^s, \\ \forall y \in Y : C_y = g^{q_y(0)}, C'_y = H(\text{att}(y))^{q_y(0)})$$

Here, $H : \{0, 1\}^* \rightarrow \mathbb{G}_1$ is a hash function, modeled as random oracle, that maps string attribute to random element of \mathbb{G}_1 .

- **KEYGEN(MK, S):** The secret key SK corresponding to a set of attributes S is (random $r, r_j \in \mathbb{Z}_p$):

$$SK = (D = g^{(\alpha+r)/\beta}, \\ \forall j \in S : D_j = g^r H(j)^{r_j}, D'_j = g^{r_j})$$

D_j, D'_j for each attribute are blinded by r_j , and all the components are tied together using r in D . This prevents attributes of different users from being combined together and provides collusion resistance.

- **DECRYPT(CT, SK):** The goal of decryption algorithm is to find out $e(g, g)^{\alpha s}$. It finds out the secret $q_x(0)$ at each node x blinded by the random value r . A secret key SK that achieves d_R such secrets at the root R , can solve the polynomial q_R and decrypt the ciphertext. A recursive algorithm DecryptNode pairs D_i and D'_i (from SK) with C_x and C'_x (from CT) respectively and returns $e(g, g)^{r q_x(0)}$ for each leaf node x in the τ in CT , iff $i = \text{attr}(x)$. $i \in S$ is the set of attributes for which a user is assigned SK .

At each non-leaf node, Lagrange interpolation is used on at least k (the threshold value of the node) such $e(g, g)^{r q_x(0)}$ received from its children z , to calculate $e(g, g)^{r q_x(0)}$. Let $A = e(g, g)^{r q_R(0)} = e(g, g)^{r s}$. Then \tilde{C}, C, D and A are used in bilinear mapping to cancel out $e(g, g)^{r s}$, and retrieve M . Further details can be found in [4].

- **DELEGATE(SK, \tilde{S}):** The delegate algorithm re-randomizes the relevant set of attributes $\tilde{S} \subseteq S$ of a secret key SK assigned for some set of attributes S . It outputs a secret key \tilde{SK} for the set of attributes \tilde{S} .

$$\tilde{SK} = (\tilde{D} = D f^{\tilde{r}},$$

$$\forall k \in \tilde{S} : \tilde{D}_k = D_k g^{\tilde{r}} H(k)^{\tilde{r}_k}, \tilde{D}'_k = D'_k g^{\tilde{r}_k})$$

Revocation Scheme of Naor and Pinkas

This scheme consists of 2 phases:

- **Initialization:** The group controller generates a random polynomial P of degree t over \mathbb{Z}_p . It sends a personal key $\langle I_u, P(I_u) \rangle$ to each user u with random identity I_u . This process is performed only once for all future revocations.
- **Revocation:** The group controller learns the random identities of t users I_{u_1}, \dots, I_{u_t} that should be revoked. At most t users can be revoked since the scheme depends on polynomial secret sharing. It then chooses a random r , and sets the new key to be $g^{r P(0)}$, that would be unknown to revoked users. It broadcasts the message $g^r, \langle I_{u_1}, g^{r P(I_{u_1})} \rangle, \dots, \langle I_{u_t}, g^{r P(I_{u_t})} \rangle$ encrypted with the current group key. Each non revoked user can compute $g^{r P(I_u)}$ and combine it with the broadcast values to obtain $g^{r P(0)}$ using Lagrange's interpolation formula. Further details can be found in [10].

3.2 Proxy-based Complete Key Revocation

In this section we describe how to completely revoke keys from parties. That means, all the privileges granted by the key authority are revoked from one or more contact(s). This construction allows revocation

of up to t users at a time since it is based on the scheme in [10] described before.

Intuition:

The master key MK contains a polynomial P of degree t . $P(0)$ is used to blind users' secret keys. Each user u also gets a random share $P(u)$ of $P(0)$ in her key. The proxy key consists of t such shares and is used to convert a part of the ciphertext for decryption. Whenever access is revoked from someone, her share becomes a part of the proxy key, and eventually the converted ciphertext. Therefore, the revoked user does not have enough points, i.e. $(t+1)$ points to unblind her key and the ciphertext and decrypt it. However, non-revoked users can always combine their secret keys with the ciphertext and hence decrypt it.

When no one is revoked, the proxy key consists of t random $P(u)$ points. Since the revocation is based on polynomial secret sharing, and the degree of the polynomial is t , the scheme is limited to maximum t revocations. Though each time t different users can be revoked, the total number of users in the system is not limited.

- **SETUP**: The key authority KA randomly generates a polynomial P of degree t (the maximum number of revoked users) over \mathbb{Z}_p , sets the broadcast secret $P(0)$ to be used after revocation, and randomly chooses $\alpha, \beta \in \mathbb{Z}_p$. She generates PK and MK as follows:

$$PK = \mathbb{G}_1, \mathbb{G}_2, g_1, g_2, h = g_1^\beta, e(g_1, g_2)^\alpha$$

$$MK = \beta, g_2^\alpha, P$$

- **ENCRYPT(PK, M, τ)**: Let Y be the set of leaf nodes in τ . Data M is encrypted to get the ciphertext CT . Other than the asymmetric groups, this algorithm works exactly the same as in BSW CP-ABE.

$$CT = (\tau, \tilde{C} = Me(g_1, g_2)^{\alpha s}, C = h^s = g_1^{\beta s},$$

$$\forall y \in Y : C_y = g_1^{q_y(0)}, C'_y = H(att(y))^{q_y(0)} = g_2^{h_y q_y(0)})$$

where $H : \{0, 1\}^* \rightarrow \mathbb{G}_2$ and $h_y = \log_{g_2} H(att(y))$ (used for notational convenience only).

- **KEYGEN(MK, S)**: The algorithm KeyGen outputs the secret key corresponding to the set of attributes S , blinded by $P(0)$ from MK . We introduce an extra component— D'_j —that in addition to attribute information contains user information. Without loss of generality, we assume user u_k receives this key.

$$SK = (D, \forall j \in S : \langle D_j, D'_j, D''_j \rangle), \text{ where}$$

$$D = g_2^{(\alpha+r)/\beta},$$

$$D_j = g_2^r H(j)^{r_j P(0)} = g_2^{r+h_j r_j P(0)},$$

$$D'_j = g_1^{r_j},$$

$$D''_j = (D'_j)^{P(u_k)} = g_1^{r_j P(u_k)}$$

- **PROXYREKEY(PK, MK, RL)**: Whenever the KA wants to revoke keys from social contacts, she creates

a list of revoked users RL with their identities u_i , $i \in \{1, \dots, t\}$, and evaluates the corresponding $P(u_i)$ using MK . She gives the proxy key PXK to the proxy. In case of no or fewer than t revocations, the KA generates random $\langle x, P(x) \rangle$ other than the actual user identities, to make PXK of length t .

$$PXK = \forall u_i \in RL : \langle u_i, P(u_i) \rangle$$

- **CONVERT(PXK, $\forall y \in Y : C_y, u_k$)**: The proxy uses its key PXK and the decryptor's identity u_k to calculate C''_y as follows:

$$\forall i, j \in \{1, \dots, t\}, k \notin \{1, \dots, t\},$$

$$\lambda_i = \frac{u_k}{u_k - u_i} \cdot \prod_{j \neq i} \frac{u_j}{(u_j - u_i)},$$

$$\forall y \in Y : C''_y = (C'_y)^{\sum_{i=1}^t \lambda_i P(u_i)} = g_2^{h_y q_y(0) \sum_{i=1}^t \lambda_i P(u_i)}$$

Since the user secret key SK is blinded by $P(0)$, she needs C''_y in addition to C_y and C'_y for decryption. The proxy also calculates λ_k and gives it to the user u_k .

- **DECRYPT(CT, SK)**: The decryption steps involve one extra pairing than BSW CP-ABE at each leaf node of the policy. For each leaf node x where $i = attr(x)$, if $i \in S$, (S is the set of attributes for which SK is issued) then,

$$\begin{aligned} DecryptNode(CT, SK, x) &= \frac{e(C_x, D_i)}{e(D''_i, C'_x)^{\lambda_k} e(D'_i, C''_x)} \\ &= \frac{e(g_1, g_2)^{r q_x(0) + h_i r_i P(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) \lambda_k P(u_k)} e(g_1, g_2)^{r_i h_i q_x(0) \sum_{j=1}^t \lambda_j P(u_j)}} \\ &= \frac{e(g_1, g_2)^{r q_x(0) + h_i r_i P(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) (\sum_{j=1}^t \lambda_j P(u_j) + \lambda_k P(u_k))}} \\ &= \frac{e(g_1, g_2)^{r q_x(0) + h_i r_i P(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) P(0)}}, k \notin \{1, 2, \dots, t\} \\ &= e(g_1, g_2)^{r q_x(0)} \end{aligned}$$

Otherwise $DecryptNode$ returns \perp . The rest of the decryption is the same as CP-ABE. For each child z of a non-leaf node x , it calculates $F_z = e(g_1, g_2)^{r q_z(0)}$. Let S_x be a threshold-sized arbitrary set of children of x , such that $F_z \neq \perp$. Then interpolation and pairings are used to calculate $e(g_1, g_2)^{\alpha s}$, and hence retrieve M .

$$F_x = \prod_{z \in S_x} F_z^{\lambda_i}, [i = index(z),$$

$$\lambda_i \text{ calculated over the indices of } z \in S_x]$$

$$= \prod_{z \in S_x} (e(g_1, g_2)^{r q_z(0)})^{\lambda_i}$$

$$= \prod_{z \in S_x} (e(g_1, g_2)^{r q_{parent(z)}(index(z))})^{\lambda_i}, [\text{discussed in 3.1}]$$

$$= \prod_{z \in S_x} (e(g_1, g_2)^{r q_x(i)})^{\lambda_i}$$

$$= e(g_1, g_2)^{\sum_{z \in S_x} r \lambda_i q_x(i)}$$

$$= e(g_1, g_2)^{r q_x(0)}$$

Let $A = e(g_1, g_2)^{rqR(0)} = e(g_1, g_2)^{rs}$ at the root R . Decryption proceeds as follows,

$$\frac{\tilde{C}}{e(C, D)/A} = Me(g_1, g_2)^{\alpha s} \frac{e(g_1, g_2)^{rs}}{e(g_1, g_2)^{\alpha s + rs}} = M$$

Explanation of Asymmetric Group:

We use different groups for C'_i and D'_i (i is the attribute). The user gets C'_i converted to $C''_i = C'^a_i$ ($a = \sum_{j=1}^t \lambda_j P(u_j)$, explained in the *Convert* algorithm) by the proxy. However, if both C'_i and D'_i belong to the same group, and the user gives the proxy D'_i instead of C'_i , she will get $(D'_i)^a = g^{ar_i}$. She will also get λ_k , which she can use to get $D''_{i\lambda_k} = g^{\lambda_k P(u_k) r_i}$. Multiplying these two, she gets $g^{r_i(a + \lambda_k P(u_k))} = g^{r_i P(0)}$. She can use this last value to decrypt any ciphertext without using the proxy, so the revocation is no longer effective. Therefore, we use asymmetric pairing, where C'_i and D'_i are in different groups and mapping of D'_i into the C'_i group is not possible.

3.3 Delegation of Access

We design delegation of attributes in PIRATTE in two settings - 1) When all the keys to different parties are issued by a single key authority, and 2) When keys are issued by multiple key authorities.

3.3.1 Single Key Authority

In the single authority setting, a user u_k gets a secret key SK from key authority KA, and delegates one or more of the attributes that she possesses in her secret key to another user. As long as u_k is not revoked, the delegated key can be used for decryption. The delegation process is as follows. The delegation algorithm takes in a secret key SK issued for a set of attributes S and delegates one or more attributes from this set. The delegated key \tilde{SK} is generated for the subset of attributes $\tilde{S} \subseteq S$. For delegation in single authority setting, an extra public parameter $f = g_2^{1/\beta}$ is introduced in the public key PK . Let random $\tilde{r} \in Z_p$, and random $\forall_j \in \tilde{S}, \tilde{r}_j \in Z_p$.

$$\begin{aligned} \tilde{SK} &= (\tilde{D}, \forall_j \in \tilde{S} : \langle \tilde{D}_j, D'_j, \tilde{D}''_j \rangle), \text{ where} \\ \tilde{D} &= (D)^{f^{\tilde{r}}} = g_2^{(\alpha + r + \tilde{r})/\beta} \\ \tilde{D}_j &= (D_j) g_2^{\tilde{r} H(j)} = g_2^{r + \tilde{r} + h_j r_j P(0) + h_j \tilde{r}_j} \\ \tilde{D}''_j &= (D''_j) g_1^{\tilde{r}_j / \lambda_k} = g_1^{r_j P(u_k) + \tilde{r}_j / \lambda_k} \end{aligned}$$

Decryption proceeds are follows:

$$\begin{aligned} \text{DecryptNode}(CT, \tilde{SK}, x) &= \frac{e(C_x, \tilde{D}_i)}{e(\tilde{D}''_i, C'_x)^{\lambda_k} e(D'_i, C''_x)} \\ &= \frac{e(g_1, g_2)^{q_x(0)(r + \tilde{r} + h_i r_i P(0) + h_i \tilde{r}_i)}}{e(g_1, g_2)^{(r_i P(u_k) + \tilde{r}_i / \lambda_k) h_i q_x(0) \lambda_k}} \cdot \frac{1}{e(g_1, g_2)^{r_i h_i q_x(0) \sum_{j=1}^t \lambda_j P(u_j)}} \\ &= \frac{e(g_1, g_2)^{(r + \tilde{r}) q_x(0) + h_i r_i P(0) q_x(0) + h_i \tilde{r}_i q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) \lambda_k P(u_k) + h_i \tilde{r}_i q_x(0)}} \cdot \frac{1}{e(g_1, g_2)^{r_i h_i q_x(0) \sum_{j=1}^t \lambda_j P(u_j)}} \\ &= \frac{e(g_1, g_2)^{(r + \tilde{r}) q_x(0) + h_i r_i P(0) q_x(0) + h_i \tilde{r}_i q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) (\sum_{j=1}^t \lambda_j P(u_j) + \lambda_k P(u_k)) + h_i \tilde{r}_i q_x(0)}} \\ &= \frac{e(g_1, g_2)^{(r + \tilde{r}) q_x(0) + h_i r_i P(0) q_x(0) + h_i \tilde{r}_i q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) P(0) + h_i \tilde{r}_i q_x(0)}}, \\ &k \notin \{1, 2, \dots, t\} \\ &= e(g_1, g_2)^{(r + \tilde{r}) q_x(0)} \end{aligned}$$

Let $A = e(g_1, g_2)^{(r + \tilde{r}) q_R(0)} = e(g_1, g_2)^{(r + \tilde{r}) s}$ for the root node. The rest of the decryption proceeds as follows,

$$\frac{\tilde{C}}{e(C, \tilde{D})/A} = Me(g_1, g_2)^{\alpha s} \frac{e(g_1, g_2)^{(r + \tilde{r}) s}}{e(g_1, g_2)^{\alpha s + rs + \tilde{r} s}} = M$$

3.3.2 Multiple Key Authority

The second version of delegation of access is designed with a distributed setting in mind, i.e., when secret attribute keys are issued from different key authorities to different users. In this setting, **A** generates keys for **B** and **B** generates keys for **C**; i.e., **B** is a contact of **A** and **C** is a contact of **B**. Again, the delegation algorithm takes in a secret key SK issued for a set of attributes S and delegates one or more attributes from this set. In the following construction, we will show how **B** delegates a key SK for the set of attributes S generated by **A** for **B**, to **C**.

$$SK = (D, \forall_j \in S : \langle D_j, D'_j, D''_j \rangle), \text{ where}$$

$$\begin{aligned} D &= g_2^{(\alpha + r)/\beta}, \\ D_j &= g_2^{r_j H(j)} g_1^{r_j P_A(0)}, D'_j = g_1^{r_j}, D''_j = (D'_j)^{P_A(B)} \end{aligned}$$

where random $r, r_j \in Z_p$, P_A is the polynomial in **A**'s MK , and B is **B**'s identity.

Attribute delegation allows **B** to delegate some subset of attributes $\tilde{S} \subseteq S$ to his contact **C**. **B** re-randomizes SK for **C** for a set of attributes $\tilde{S} \subseteq S$.

$$\begin{aligned} \tilde{SK} &= (D, \forall_j \in \tilde{S} : \langle D_j, D'_j, \tilde{D}''_j, \tilde{D}'''_j \rangle), \text{ where} \\ \tilde{D}''_j &= (D''_j)^{1/P_B(0)}, \tilde{D}'''_j = (D''_j)^{P_B(C)/P_B(0)} \end{aligned}$$

where P_B is the polynomial in **B**'s MK , and C is **C**'s identity.

To decrypt a ciphertext CT encrypted with A 's public parameters, A 's proxy calculates $C''_{yA} = (C'_y)^{X_A}$ and λ_B , B 's proxy calculates $C''_{yB} = (C'_y)^{X_B}$ and λ_C , and gives it to C . Here, C'_y is from CT , and X_A and X_B are revocation information for A 's proxy and B 's proxy respectively, calculated as $X_A = \sum_{i=1}^{t_1} \lambda_i P_A(u_i)$, $X_B = \sum_{j=1}^{t_2} \lambda_j P_B(u_j)$. u_i and u_j are the revoked users by A and B respectively, λ_i and λ_j are the Lagrange's coefficients for corresponding revoked users, and t_1 and t_2 are the parameters for the maximum number of revoked users by A and B defined in their MK s. λ_B and λ_C are the Lagrange's coefficients for B and C respectively. B 's identity is conveyed to C as a part of the communication for \tilde{SK} , or any other way. C does a modified bilinear pairing in the *DecryptNode*, and finally decrypts the data.

$$\begin{aligned}
& DecryptNode(CT, \tilde{SK}, x) \\
&= \frac{e(C_x, D_i)}{e(\tilde{D}_i'', C''_{xB})^{\lambda_B} e(\tilde{D}_i''', C'_x)^{\lambda_B \lambda_C} e(D'_i, C''_{xA})} \\
&= \frac{e(C_x, D_i)}{e(g_0, g_1)^{r_i h_x q_x(0) X_B \lambda_B P_A(B) / P_B(0)}} \\
&\quad \cdot \frac{1}{e(g_0, g_1)^{r_i h_x q_x(0) P_A(B) P_B(C) / P_B(0) \lambda_B \lambda_C} e(D'_i, C''_{xA})} \\
&= \frac{e(g_0, g_1)^{r q_x(0) + r_i h_x q_x(0) P_A(0)}}{e(g_0, g_1)^{r_i h_x q_x(0) \lambda_B P_A(B)} e(g_0, g_1)^{r_i h_x q_x(0) X_A}} \\
&= \frac{e(g_0, g_1)^{r q_x(0) + r_i h_x q_x(0) P_A(0)}}{e(g_0, g_1)^{r_i h_x q_x(0) P_A(0)}} \\
&= e(g_0, g_1)^{r q_x(0)}
\end{aligned}$$

The rest of the decryption proceeds as before. If A revokes B or B revokes C , the decryption will not succeed since both the proxies participate in the decryption, and the delegated secret key \tilde{SK} contains information about all A , B , and C .

3.4 Proxy-based Attribute Revocation

In this section, we describe how to revoke one or more attributes from a given secret key. This is useful since often the KA may want to merely revoke a few attributes from her contacts instead of the whole key. For instance, user A might want to remove friend attribute from B , but B still remains in her colleague group.

Intuition:

The idea is basically the same as complete key revocation. The master key contains one polynomial P_i of degree t_i for each possible attribute i that the KA can assign. Any attribute can be introduced later by introducing a new polynomial in the MK . $P_i(0)$ is used to blind the corresponding attribute in the secret keys. Each user u also gets a random share $P_i(u)$ of $P_i(0)$ in her key. The proxy key consists of t_i such shares for each attribute in the policy used in the ciphertext. Whenever some attribute is revoked from

some user, that share becomes a part of the proxy key, and hence the converted ciphertext. Therefore, the revoked user does not have enough points, i.e. $(t_i + 1)$ points for that specific attribute to unblind her key and the ciphertext and decrypt it. However, non-revoked users can always combine their secret keys with the ciphertext and hence decrypt it. As before, when no attribute is revoked, the proxy key consists of t_i random points for each attribute i .

- **Setup:** The KA generates one polynomial P_y randomly over \mathbb{Z}_p for each attribute $y \in Y'$ where Y' is an initial set of attributes in the system, and sets $P_y(0)$ as the secret to be used to revoke the attribute. To revoke an attribute from t users at a time, the degree of the polynomials is chosen to be t . New attributes can be introduced later by randomly generating polynomials for them. Finally, she randomly chooses $\alpha, \beta \in \mathbb{Z}_p$.

$$\begin{aligned}
PK &= \mathbb{G}_1, \mathbb{G}_2, g_1, g_2, h = g_1^\beta, e(g_1, g_2)^\alpha \\
MK &= \beta, g_2^\alpha, \forall y \in Y' : P_y
\end{aligned}$$

- **KeyGen(MK, S):** The components of the secret key are similar as before except that the polynomial in each is specific to the attribute represented by the component. Again, without loss of generality, we assume user u_k receives this key.

$SK = (D, \forall j \in S : \langle D_j, D'_j, D''_j \rangle)$, where

$$\begin{aligned}
D &= g_2^{(\alpha+r)/\beta}, D_j = g_2^{r_j} \cdot H(j)^{r_j P_j(0)} = g_2^{r+h_j r_j P_j(0)}, \\
D'_j &= g_1^{r_j}, D''_j = (D'_j)^{P_j(u_k)} = g_1^{r_j P_j(u_k)}
\end{aligned}$$

- **Encrypt(PK, M, τ):** Encryption is similar as in complete key revocation.

- **ProxyRekey(PK, MK, $\forall y \in Y : RL_y$):** To revoke an attribute $y \in Y$ from t contacts, the KA creates a t -sized list $RL_y = \{u_i\}$, $i \in \{1, \dots, t\}$ of revoked users for that attribute, and evaluates $P_y(u_i)$ using MK . In case of no or less than t revocations, she generates random $\langle x, P_y(x) \rangle$ to make RL_y of length t . The set of users from whom different attributes are revoked, may or may not overlap. Without loss of generality we assume that the sets of revoked users don't overlap. The proxy key PXK is constructed as follows:

$$PXK = \forall y \in Y, \forall u_i \in RL_y : \langle u_i, P_y(u_i) \rangle$$

- **Convert(PXK, $\forall y \in Y : C_y$):** The proxy uses its key PXK to convert the attribute relevant components C'_y received from user u_k to C''_y as follows:

$$\begin{aligned}
\lambda_i^y &= \frac{u_k}{u_k - u_i} \cdot \prod_{j \neq i} \frac{u_j}{(u_j - u_i)}, \\
&\forall u_i, u_j \in RL_y, u_k \notin RL_y, RL_y \in PXK
\end{aligned}$$

$$\forall y \in Y : C''_y = (C'_y)^{\sum_{i=1}^t \lambda_i^y P_y(u_i)} = g_2^{h_y q_y(0) \sum_{i=1}^t \lambda_i^y P_y(u_i)}$$

$\forall y \in Y$ the proxy also calculates and gives $\lambda_{j,k}^y$ to u_k .

- **Decrypt(CT, SK):** For each leaf node x where $i = \text{attr}(x)$, if $i \in S$ (S is the set of attributes for which SK is issued), and i is not revoked from u_k then,

$$\begin{aligned}
& \text{DecryptNode}(\text{CT}, \text{SK}, x) \\
&= \frac{e(C_x, D_i)}{e(D_i'', C'_x)^{\lambda_k^i} e(D_i', C''_x)} \\
&= \frac{e(g_1, g_2)^{r_{q_x}(0) + h_i r_i P_i(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) \lambda_k^i P_i(u_k)} e(g_1, g_2)^{r_i h_i q_x(0) \sum_{j=1}^t \lambda_j^i P_i(u_j)}} \\
&= \frac{e(g_1, g_2)^{r_{q_x}(0) + h_i r_i P_i(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) (\lambda_k^i P_i(u_k) + \sum_{j=1}^t \lambda_j^i P_i(u_j))}} \\
&= \frac{e(g_1, g_2)^{r_{q_x}(0) + h_i r_i P_i(0) q_x(0)}}{e(g_1, g_2)^{r_i h_i q_x(0) P_i(0)}}, k \notin \{1, 2, \dots, t\} \\
&= e(g_1, g_2)^{r_{q_x}(0)}
\end{aligned}$$

Otherwise DecryptNode returns \perp . The rest of the decryption is as before. In summary, if an attribute i is revoked from user u , he can not do pairing on C'_x and D'_i . He can continue to use components related to his other unrevoked attributes. Therefore, some of his attributes are revoked whereas some continue to be active.

4 SECURITY ANALYSIS

First, we need to define the requisite security properties for CP-ABE with Proxy Revocation. We present the definition for identity-based revocation; the definition for attribute-based revocation is analogous. We base our definition on the security model defined by Bethencourt et al. [4], with the addition of revocation and proxy operations. In this game, all encryptions remain secure even when the adversary compromises the proxy and obtains its key material, as long as this happens after the most recent revocation.

Setup. The challenger runs the SETUP algorithm and gives the public parameters, PK , to the adversary. The challenger also runs $\text{PROXYREKEY}(PK, MK, \emptyset)$ to generate a proxy key PXK .

Phase 1. The adversary makes repeated queries to KEYGEN to obtain keys for users u_1, \dots, u_{q_1} with sets of attributes S_1, \dots, S_{q_1} . The adversary also interacts with the proxy by calling CONVERT with the input $(\{C'_1, \dots, C'_r\}, u_k)$ for $C'_i \in \mathbb{G}_1$ and $u_k \in \mathbb{Z}_p$, at which point the challenger runs the CONVERT algorithm with the stored proxy key PXK . Finally, the adversary may call PROXYREKEY by supplying a revocation list RL . This will cause the challenger to update the proxy key PXK .

Challenge. The adversary submits two equal length messages M_0 and M_1 and an access structure \mathbb{A}^* . The adversary also supplies a new revocation list RL^* . RL^* and \mathbb{A}^* satisfy the constraint that, for each user u_k , either $u_k \in RL^*$ or S_k does not satisfy \mathbb{A}^* .

The challenger flips a coin to obtain a random bit b and returns M_b encrypted with the access structure

\mathbb{A}^* . Additionally, it runs $\text{PROXYREKEY}(PK, MK, RL)$ and returns the resulting key PXK to the adversary.

Phase 2. The adversary makes repeated queries to KEYGEN to obtain keys for users $u_{q_1+1}, \dots, u_{q_2}$ with attribute sets $S_{q_1+1}, \dots, S_{q_2}$. The new keys have to satisfy that if $u_k \notin RL^*$, then S_k does not satisfy \mathbb{A}^* .

Guess. The adversary outputs a guess b' of b .

The advantage of an adversary is defined as $\Pr[b' = b] - \frac{1}{2}$.

Definition 1: A ciphertext-policy attribute-based encryption with proxy revocation scheme is secure if all polynomial time adversaries have at most negligible advantage in the above game.

4.1 Proof Sketch

We can prove the security of our scheme using a variant of a generic bilinear group model. Note that since the security of the original CP-ABE scheme relies on the generic bilinear group model, the assumption we make is only slightly stronger than the original. In particular, we must work within an asymmetric bilinear group, with a pairing of $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, such that there is no efficiently computable isomorphism from \mathbb{G}_1 to \mathbb{G}_2 . (In a symmetric bilinear group, a user could submit D'_j to CONVERT and recover $g_0^{r_j P(0)}$, obviating the need to use the proxy in further decryptions.) This is believed to hold true for MNT curves [12].

The generic asymmetric bilinear group model. Consider three random encodings of the additive group \mathbb{F}_p represented by injective maps $\psi_1, \psi_2, \psi_T : \mathbb{F}_p \rightarrow \{0, 1\}^m$, where $m > 3 \log p$. We will define $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ as the range of the respective map. We are given access to a group action oracle for each group and an oracle for a non-degenerate bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ (we will refer to the ranges of ψ_1, ψ_2, ψ_T as $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$, respectively). We are also given oracle access to the isomorphism $\phi : \mathbb{G}_2 \rightarrow \mathbb{G}_1$ and a hash function $H : \{0, 1\}^* \rightarrow \mathbb{G}_2$. Finally, we let $g_i = \psi_i(1)$ for $i = 1, 2$.

Theorem 1: The construction presented in Section 3 is secure under the generic asymmetric bilinear group model.

We sketch the main argument here; most of the rest of the details are similar to the proof presented by Bethencourt et al. [4].

Sketch: First of all, we can assume that no “unexpected collisions” happen between the maps, meaning that if we keep track of the algebraic expressions passed to ψ_1, ψ_2, ψ_T , and ϕ , two values are equal if and only if the expressions are symbolically equivalent. This assumption is true except for a negligible probability.

For simplicity of presentation, we will assume that \mathbb{A}^* contains a single attribute A_j for some j . Then after phase 2, the adversary has the following elements available:

$\mathbb{G}_1 : g_1, g_1^\beta, C = g_1^{\beta s}, C_j = g_1^s (= g_1^{q_j(0)}), C'_j g_1^{h_j s}$, where s is the random secret used to encrypt the challenge message and h_j is implicitly defined such that $H(j) = g_2^{h_j}$.

$$\mathbb{G}_2 : g_2, D = g_2^{(\alpha + r_{u_k})/\beta}, D_j = g_2^{r_u + h_j r_{u_k, j} P(0)}, \\ D_j = g_2^{r_{u_k, j}}, D'_j = g_2^{r_{u_k, j} P(u_k)}$$

for each queried user u_k where $A_j \in S_k$. Note that we can ignore all $D_{j'}$ for $j' \neq j$ because, as with CP-ABE, they will not help with decryption.

$$\mathbb{G}_T : e(g_1, g_2)^\alpha, M \cdot e(g_1, g_2)^{\alpha s}$$

In addition, the adversary knows $u_k, P(u_k)$ for all the revoked users in RL^* . Note that we can ignore any other elements of \mathbb{G}_1 obtained through calls to CONVERT during Phase 1, or through calls to the isomorphism. This is because in order to guess correctly with a non-negligible probability, the adversary needs to compute $e(g_1, g_2)^{\alpha s}$. Since there are no occurrences of s in \mathbb{G}_2 , each pairing must involve g_1^{sk} for some k , and hence be derived from C, C_j , or C'_j .

The adversary can compute $e(C, D^{(u)}) = e(g_1, g_2)^{\alpha s + r_u s}$, hence computing $e(g_1, g_2)^{\alpha s}$ is equivalent to computing $e(g_1, g_2)^{r_u s}$ for some user u . Note also that the secret keys obtained for other users are *not* helpful here, for the same reason as in original CP-ABE. Algebraically, the adversary must solve the following equation:

$$e(g_1, g_2)^{r_u s} = e\left(g_1^x, \left(D_j^{(u)}\right)^y\right) e(g_1, g_2)^z$$

where x, y , and z are derived from the available elements *other* than $D_j^{(u)}$. Note that $x, y \neq 0$, since otherwise there is no way to introduce r_u into the right-hand side. However, the rest of the values are *independent* of $P(0)$, since the only values of the polynomial available to the adversary outside of $D_j^{(u)}$ are $P(u_k)$ for $u_k \in RL^*$, and thus are not sufficient to determine $P(0)$. \square

5 CASE STUDY: SOCIAL NETWORKS

Online Social Networks (OSNs) such as Facebook, Google+, Twitter, and LinkedIn are becoming one of the most popular ways for users to interact online. Besides personal communication, OSNs provide the perfect platform for online games and other applications. Users share personal information with the social network provider, and trust the provider to protect their sensitive information. However, this introduces privacy risks, as the collection of information is an attractive attack target [13]. Insiders can also release private information either intentionally or accidentally [14], [15]. Several recent privacy compromises have thrown these issues into sharp focus [1], [2].

These issues have motivated researchers to consider a paradigm shift, where instead of trusting social network operators and being dependent on them to enforce privacy, *users* are in control of who views their data, for example, via encryption [7], [8], [9], [16]. Fine-grained access control is a key challenge in this space; for example, Facebook and LiveJournal have rolled out mechanisms to specify access control policies for each post, as the data items are usually destined for a subset of friends, or groups.

Persona [7] is a state-of-the-art design that proposes the use of ABE to enable fine-grained access control. In OSNs, ABE allows users to have complete control over who can see their data, free from the whims of the OSN provider. A user can create groups by assigning different attributes and keys to her social contacts, and then encrypt data such that only particular users having the desired set of attributes can decrypt it. This provides information protection from unauthorized users on the OSN, third-party application developers, and above all the OSN provider itself.

However, groups are dynamic and therefore user attributes may change over time. This could be because of change in location, work environment, or the nature or strength of the relationship with a contact. Recent studies have shown that the user interaction graph is much less dense than friendship graph [17], indicating that users interact most frequently with a small group of friends, further validating the need for fine-grained access control. Moreover, the churn rate for the interaction graph has been shown to be quite high [17], motivating the need for access control mechanisms to support *dynamic* groups.

Persona and similar designs introduce significant overhead for group membership changes, especially when a contact is removed from a group: all other members of the group must receive a new key; additionally, all existing data items destined for that group must be re-encrypted. This does not scale when group sizes are large and group churn rate is high.

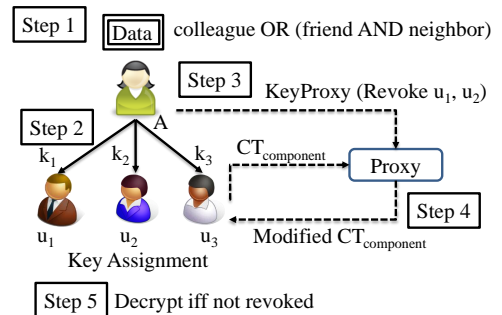


Fig. 1. Architecture of EASIER- OSN Built Using PIRATTE

We propose an architecture for OSN using PIRATTE as the underlying cryptographic scheme. Encryption

successfully hides information from unintended parties. Traditional OSNs allow people to establish only a common relationship with each other by adding them as friends. Further specialization is achieved by creating lists and adding friends to that list. Relationships based on attributes make this task more expressive.

Figure 1 shows the architectural overview of an OSN built using PIRATTE. We call this architecture EASiER [18]. Users in EASiER become their own key authorities, define relationships by assigning attributes and relevant keys to the parties, and encrypting data under a policy that, if satisfied, allows decryption to the intended parties. For instance, user A assigns keys for attributes (*colleague*, *neighbor*) to user B and encrypts data for the policy '*colleague OR (friend AND neighbor)*'.

As mentioned earlier, the user interaction graph is different from the friendship graph. People interact with a subset of their defined contacts most of the time. Recent privacy setting changes in wall posts in Facebook also supports this fact. This requirement needs ABE in OSNs to consider revocation of access without re-encrypting all the data and re-keying everyone in the group when access is denied to a contact. For example, after encrypting data under the mentioned policy, user A might want everyone except u_1 and u_2 in her *colleague* group to decrypt the ciphertext, either temporarily or permanently. Besides, A might want to revoke the *colleague* attribute from some of her contacts. PIRATTE provides this option by introducing a minimally trusted proxy.

Upon revocation, the owner supplies enough information, in this case, a set of t (the maximum number of revoked users allowed) revoked users to the proxy to construct a proxy key. The proxy is minimally trusted. Hence, the OSN provider or a third party can act as proxy. An unrevoked user sends a specific set of components of the ciphertext (details described earlier) to the proxy. The proxy uses its key to change these components such that only unrevoked users can use it, mathematically combine it with the rest of the ciphertext and their keys, and finally obtain the plaintext. Since the conversion involves only the lightweight components of ciphertext, it is not expensive, as we will demonstrate later through experiments.

However, PIRATTE does not allow the proxy to decrypt the data since it does not have the attribute keys. A new proxy key, created each time a revocation takes place, prevents revoked users to collude with the proxy or with each other to get the data. This prevents the revoked users from decrypting even old data unless they store it somewhere. We argue that this is a desirable property: currently trusted contacts are not likely to crawl the entire set of social network data and store it for later use, but former friends or colleagues might try to abuse their former status by accessing past data.

Another example of an OSN that utilizes PIRATTE is DECENT, a decentralized architecture for OSN [19]. It uses PIRATTE for cryptographic protection, DHT for efficient data storage, and an object oriented architecture for flexible data representation.

Friend-of-friend:

A challenge in using cryptography to enforce access control in OSN is to support degrees of relationships, such as *friend-of-friend* or *contact-of-contact* to be more generic. EASiER handles this challenge by using the feature attribute delegation in the distributed setting described before.

In this approach, a user A generates a secret key for her contact B with the attributes she wants to assign to him. She also adds an attribute named *fof* to the key. B uses the access delegation algorithm to delegate this specific attribute (*fof*) to his contacts, for example C. Whenever A encrypts a piece of data with the policy *attr1, attr2, ... OR fof*, C can decrypt it with the delegated key that he received from B. As in distributed attribute delegation, if A revokes B, or B revokes C, the decryption will not succeed since both the proxies are required to participate in the decryption, and the delegated secret key SK contains information about A, B, and C.

6 IMPLEMENTATION AND EXPERIMENTAL EVALUATION

We implemented the constructions in PIRATTE, as described in Section 3. Our implementation involves introducing new components as well as modifying different parts of the BSW CP-ABE toolkit [20]. The current implementation supports complete key revocation and access delegation in a distributed setting. Similar techniques can be applied to modify it to perform attribute revocation.

The implementation uses MNT curves [12] with 159 bit base field. All the experiments were carried out on a 2.40 GHz Intel Core 2 Duo, 4GB memory, and running Ubuntu 10.04. We also implemented a Facebook application to provide the functionality on a social network. The code and the Facebook application are available at <http://www.soniajahid.com> and <http://apps.facebook.com/myeasier> respectively.

6.1 Performance Analysis

We provide some information on the performance evaluation of PIRATTE, and compare it with CP-ABE both with MNT and super-singular curves. Though CP-ABE implementation uses symmetric pairing, we use asymmetric pairing for both PIRATTE and CP-ABE in our implementation. This provides security by preventing key and ciphertext components exchange (discussed in Section 3). The results are shown in Figure 2.

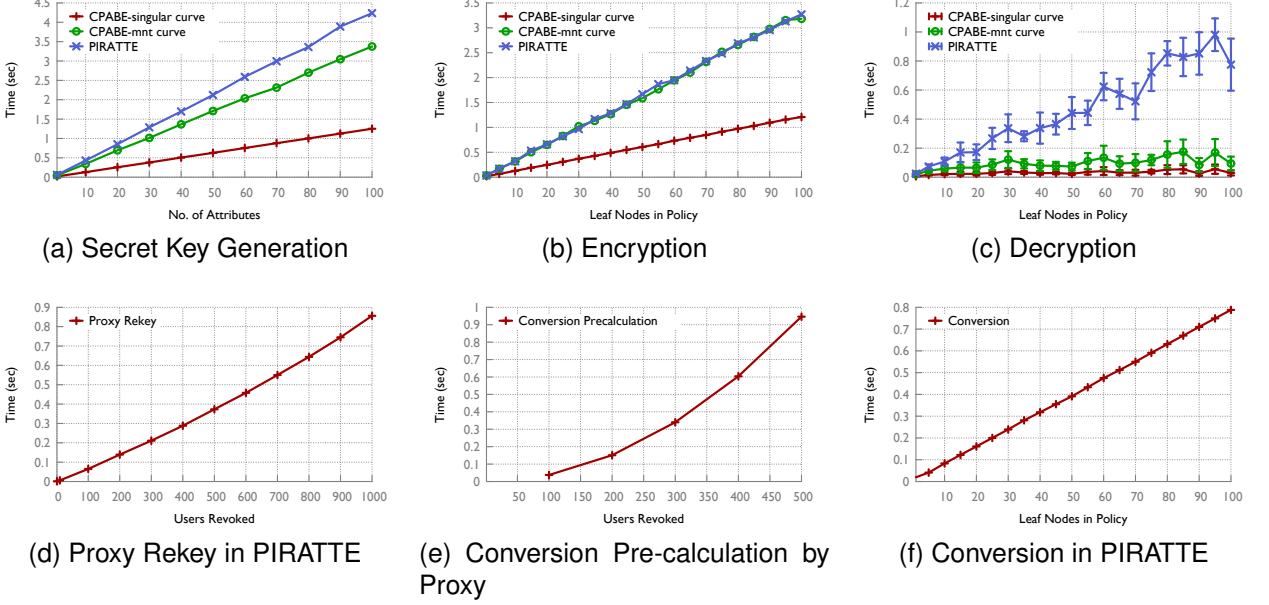


Fig. 2. Performance Analysis of PIRATTE and Comparison with BSW CP-ABE

Key Generation: Key generation time is linear with number of attributes both in CP-ABE and PIRATTE. Since it does an extra exponentiation, and generates an extra component for each attribute in PIRATTE, the result is justified. CP-ABE with supersingular curve requires the least time.

Encryption: To test encryption and decryption, we randomly generated 10 different policies for each of the desired number of leaves (1, 5, 10, ... 100). Encryption is also linear with respect to the number of leaf nodes in the policy. We did not make any changes to the encryption scheme, so CP-ABE with MNT and PIRATTE both take the same amount of time.

Decryption: Decryption depends on the policy used in encryption, and the attributes involved. We generated a decryption key with 100 attributes. It has the superset of all the attributes used to generate the policies, and so satisfies each of them. The decryption results are shown with a 95% confidence interval. All the lines show the performance when an optimization was used to ensure the usage of the minimum number of leaves in the algorithm *DecryptNode*. The required time is still below 1 second.

Proxy Rekey and Conversion: PIRATTE involves two extra costs before decryption: re-keying the proxy and converting the ciphertext components specific to the leaves in the policy. The re-keying results (Figure 2d) show that for even 1000 revoked users, the time required is less than 0.9 seconds. This should be compared with the time required to rekey the rest of a group, i.e., generate a new key for everyone, when even one person in the group is revoked.

Conversion primarily involves one exponentiation for each of the leaf specific ciphertext components. It

also calculates λ_k for the requester u_k , and completes the λ_i s for each of the revoked users. We perform an optimization by allowing the proxy to pre-calculate a portion of the λ_i 's in *Convert*. With the optimization, the proxy needs to do 1 multiplication per revoked user to calculate λ_i . It works as follows:

$$\lambda'_i = \prod_{u_i, u_j \in RL, i \neq j} \frac{u_j}{(u_j - u_i)}, \text{ and } l'_i = \lambda'_i P(u_i)$$

$$l_i = l'_i \frac{u_k}{(u_k - u_i)} = \lambda'_i \frac{u_k}{(u_k - u_i)} P(u_i) = \lambda_i P(u_i),$$

$$\forall u_i \in RL, u_k \notin RL$$

Figure 2e shows the time required for the proxy to do the pre-calculation. Note that this is a one-time cost each time the proxy is re-keyed and is not faced by users. Figure 2f shows the time required to actually convert a ciphertext. The results are almost equal for the number of revoked users since time to do t exponentiation dominates the time to do t multiplication. Figure 2f shows the time for 500 revoked users. We expect the proxy to be more powerful in terms of computing, and hence rekeying, and conversion should be faster in practice. A user performing decryption only faces the conversion time shown in Figure 2f along with the decryption time mentioned earlier.

Decryption with Delegated Key: We measured the time required to decrypt a ciphertext with delegated secret key in the multiple authority setting since we used it in EASiER. We are interested in just one attribute, i.e., *fof* being delegated for the OSN setting. The policies used to encrypt the data contain the *fof* attribute in the form: (*attr1*, ...) OR *fof*. Therefore, the

decryption needs to verify one attribute in the policy. Hence, the time required does not depend on the number of leaf nodes in the policy, and is constant. In our experiment, this time is about 0.34sec.

TABLE 1
Element Size

Group	Size (bytes)
\mathbb{G}_1	44
\mathbb{G}_2	124
\mathbb{G}_T	124
Z_p	24

TABLE 2
Component Size

Component	PIRATTE (bytes)	CP-ABE (bytes)
Public Key	1316	1316
Master Key	$152 + (t + 1)24$	148
Private Key	$128 + (a + 212)n$	$128 + (a + 168)n$
Ciphertext	$168 + 8i + (176 + a)l$	$168 + 8i + (176 + a)l$
Proxy Key	$24t$	NA
C''_y	$124l$	NA

Component Size and Communication Overhead:

Table 2 shows the sizes of the components involved in the system for complete key revocation. Size of the components for attribute revocation can be calculated similarly. Elements from $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$, and Z_p require 44, 124, 124, and 24 bytes respectively to represent (Table 1). Users have to communicate with the proxy for conversion by sending C''_y , and receiving C''_y . These are represented using elements from \mathbb{G}_2 . This requires 124 bytes to represent (120 for the actual data, and 4 for the variable size). Hence, conversion of a ciphertext with l leaf nodes in the policy will need to transfer $124l$ bytes each way. The user also sends u_k , and receives λ_k back. These are represented using Z_p which requires 24 bytes.

Public Key consists of a string describing the pairing used (980 bytes), g_1 and h from \mathbb{G}_1 , g_2 from \mathbb{G}_2 , and $e(g_1, g_2)^\alpha$ from \mathbb{G}_T . Master Key consists of β from Z_p , and g_2^α from \mathbb{G}_2 in CP-ABE. In PIRATTE, it also consists of a polynomial of degree t . The polynomial consists of an integer t , and $t + 1$ coefficients from Z_p . Private Key consists of D from \mathbb{G}_2 , number of attributes n (integer), and n of $\langle D_j, D'_j \rangle$ s from \mathbb{G}_2 and \mathbb{G}_1 respectively and attributes of length a (string). PIRATTE also contains n of D''_j from \mathbb{G}_1 . Ciphertext consists of \tilde{C} from \mathbb{G}_T , C from \mathbb{G}_1 , and components for each node in the policy. Both intermediate (i nodes) and leaf (l nodes) nodes have a threshold k (integer), and number of children (0 for leaf, also an integer). A leaf node has a string attribute of length a , and C_y and C''_y from \mathbb{G}_1 and \mathbb{G}_2 respectively. Proxy Key consists of t Z_p elements.

Attribute Revocation: We can estimate the time required to perform attribute revocation. All the algorithms work similarly. The only extra work is in Setup where instead of just one polynomial, a polynomial is generated for each attribute in the system.

It requires 0.08sec to generate a polynomial of degree 100 and increases linearly with the degree of the polynomial. Hence, generating a polynomial for each attribute is scalable.

6.2 Facebook Application

Finally, we developed a Facebook application for PIRATTE as a proof of concept. The goal is to present a high level idea of how an OSN that uses PIRATTE will look like. We focused on Facebook because of lack of deployment of decentralized architectures like Persona and Diaspora [21]. In the current implementation, all protocol functions are performed at our application server. Moreover for convenience, we chose the client's proxy server to be the application server itself.

The application retrieves public profile information of the users installing it using Facebook API, and uses this data as its user profile information. Figures 3a and 3b depict screen-shots of the same. Data is hidden for privacy purposes.

When a user first installs our Facebook application, an account is created in the application server. We provide a brief description of the supported functionality.

- **Setup:** The *Setup* button is used to initialize a public key PK and a master secret key MK.
- **Key Proxy:** The *Key Proxy* button is used to generate a key for the proxy server. This is used to key the proxy when there is no revocation. Each revocation updates the proxy key, so there is no need to manually rekey the proxy with each revocation.
- **Add Attributes:** The *Add Attributes* functionality is invoked to define a set of attributes like *family*, *coworker*, *researcher*, etc.
- **KeyGen:** The *KeyGen* functionality is used to generate keys for Facebook friends who also have the PIRATTE application installed. These users are assigned specific attributes, with the selection of attributes being made from amongst the choices defined in the Add Attributes step.
- **Encrypt:** The *Encrypt* button is used to generate a ciphertext under a policy defined over the available set of attributes. For convenience, application users can choose and encrypt userid, about, and gender for their contacts which is already retrieved from their profile information (if available).
- **Revoke:** Users can select one or more contacts to whom they assigned keys previously, and revoke the keys from them.
- **Decrypt:** A user can click on any ciphertext generated by a contact who assigned her a secret key, and is able to view the plaintext if his/her attributes satisfy the ciphertext-policy. When a user clicks Decrypt, the ciphertext is converted by the proxy, and then the user secret key is used to decrypt the

data. Names are not shown in the snapshot because of privacy.

- **Delegate:** The Delegate option is used to delegate the *fof* attribute from the keys a user received to her contacts. It shows the list of contacts from whom a user received secret attribute keys and to whom she can delegate access to. Duplicates and self-assignment are prohibited through checks.

(a) Main Page

(b) Revoke

Fig. 3. PIRATTE on Facebook

7 RELATED WORK

7.1 Revocation Schemes

Attrapadung and Imai address the revocation problem in [22] by combining broadcast encryption scheme with both CP-ABE and KP-ABE. This scheme requires knowledge about the list of all possible users during encryption, i.e., attaches one component per user to the ciphertext. Knowing the list of all possible users in advance is inconvenient for most scenarios, specifically OSNs. Bethencourt et al. in [4] and Boldyreva et al. in [23] propose expiration-based revocation scheme for CP-ABE and KP-ABE respectively. In these schemes, the secret keys (CP-ABE) or the ciphertext (KP-ABE) contain expiry time as an extra

attribute. Time-based revocation may not be a desired property in several applications where an immediate revocation is necessary. It introduces a window of vulnerability, i.e., the gap between the desired time from revocation to the actual time of revocation. Ostrovsky et al. [24] present a new KP-ABE scheme with non-monotonic access structure to support negation of attributes. Revocation can be implemented by adding a NOT-attribute to the policy in private key. However, this will require re-keying the users from whom an attribute is revoked.

Lewko et al. in [25] propose a revocation scheme with small private keys. However, their approach also increases the size of the ciphertext by incorporating the list of revoked and unrevoked users with it. In [26] Hur proposes a revocation scheme for CP-ABE where each leaf node, i.e. attribute in the policy used to encrypt the data contains a group of users who possess that attribute. The scheme consists of a key generation center (KGC) and a data storing center (DSC). The DSC is equivalent to the proxy in our proposed scheme. However, the approach requires secret key update for all the users of an attribute group from which at least one user was revoked; this also includes the non-revoked users in that group.

The CCA secure construction of Yu et al. [6] involves updating the master key component of each attribute that has been revoked in the system. The public key components are then updated and data is encrypted with the new public key. Finally, proxy re-keys are generated that enable a proxy to update the user secret keys to the new version for all but the revoked user. Basically the re-keying burden is placed on the proxy, instead of users. We note that the existing data would need to be re-encrypted by the proxy; placing a significant burden on the system. While our scheme is only CPA secure under the generic group model (weaker security than Yu et al.), it does not require re-encryption of existing data.

7.2 Proxy Based Re-encryption

Our revocation techniques are based on the notion of proxy re-encryption. We will now briefly trace some developments in this field, and discuss why the state of the art techniques cannot be directly applied to our setting.

Blaze et al. [27] introduced the notion of *proxy re-encryption*, in which a proxy could convert a ciphertext for Alice into a ciphertext for Bob, using a specially generated proxy key. The holders of public-key pairs A (Alice) and B (Bob) create and publish a proxy key $\pi_{A \rightarrow B}$ such that $D(\pi(E(m, e_A), \pi_{A \rightarrow B}), d_B) = m$, where $E(m, e)$ is the public encryption function of message m under encryption key e , $D(c, d)$ is the decryption function of ciphertext c under decryption key d , $\pi(c, \pi_{A \rightarrow B})$ is the proxy function that converts ciphertext c according to proxy key $\pi_{A \rightarrow B}$, and e_A, e_B ,

d_A, d_B are the public encryption and secret decryption component keys for key pairs A and B , respectively. In the El-Gamal based construction proposed by Blaze et al., while the proxy cannot see the plaintext message m , it can collude with B to recover the secret key for A . Moreover, the construction is *bidirectional*, and the proxy key can be used to convert ciphertext for Bob into a ciphertext for Alice as well.

Ateniese et al. [28] propose *unidirectional* protocols for proxy re-encryption based on bilinear maps, where a re-encryption key from A to B does not imply a re-encryption key from B to A . Canetti and Hohenberger [29] proposed the first CCA secure bidirectional proxy re-encryption scheme, while Libert and Vergnaud [30] proposed the first CCA secure unidirectional proxy re-encryption scheme in the standard model. We note that all of the above schemes are limited to the public key setting. Green and Ateniese [31] extended the model to the identity based setting by proposing a scheme for identity based proxy re-encryption, but it was not until Liang et al. [32] that the attribute based encryption setting was considered.

In the attribute based setting of Liang et al., a user could designate a proxy, who can re-encrypt a ciphertext with a certain access policy into another ciphertext with a different access policy. Furthermore, the authors showed their scheme to be selective-structure chosen plaintext secure. However, it is not possible to apply their construction to the problem of attribute revocation because their construction does not support negative attributes.

7.3 Social Network Privacy Architectures

We will now describe existing social network privacy architectures and provide a comparison with EASiER. We can categorize the different schemes based on the level of trust for centralized services: a) trusted central servers, b) untrusted central servers, and c) decentralized architecture. Intuitively, decentralized architectures provide the highest level of privacy. EASiER is close to Persona [7], which is the state-of-the-art decentralized architecture for social network privacy. In almost all of the settings, access control is performed via encryption. We believe that the public key encryption setting is not well suited for fine-grained access control. Instead, PIRATTE uses attribute based encryption techniques to enable users to perform fine-grained access control. Moreover, none of the schemes focus on the issue of efficient user/attribute revocation.

Trusted Centralized Architectures: Lucas and Borisov [9] propose flyByNight, a facebook application designed to mitigate privacy risks in social networks. flyByNight users encrypt sensitive messages using JavaScript on the client side and send the ciphertext to some intended party, i.e., Facebook friends, who can then decrypt the data. The architecture ensures that transferred sensitive data cannot be

viewed by the Facebook servers in an unencrypted form. However, the utility of flyByNight is limited to preserving the privacy of messages intended for social network friends, i.e., email type communication, and thus, it does not provide complete privacy. For example, the application server knows a user's friendlist on facebook. flyByNight is also vulnerable to active attacks by the OSN provider, since the OSN interface is used for key management.

Singh et al. [33] propose the xBook framework for building privacy preserving social network applications. xBook uses information flow models to control what untrusted applications can do with the information they receive. Their design retains the functionality offered by existing online social networks. xBook provides enforcement for both user-user access control for data flowing within a single application, as well as for information sharing with entities outside xBook. Social network applications are re-designed to have access to all the data that they require, but this data is not allowed to be passed on to an external entity unless approved by the user.

Untrusted Centralized Architectures: Guha et al. [8] propose to improve user privacy while still preserving the functionality of existing online social network providers. Their architecture is called *None of Your Business* (NOYB), in which encryption is used to hide the data from the untrusted social network provider. The key feature of their architecture is a general cipher and encoding scheme that preserves the semantic properties of data such that it can be processed by the social network provider *oblivious* to encryption. A user's private information is partitioned into *atoms*, and NOYB encrypts a user's atom by substituting it with the atom of another user. Thus the OSN can operate on ciphered data, but only the authorized users can decrypt the result.

Anderson et al. [34] propose a client-server architecture for providing social network privacy. In their design, the server is a very simple untrusted social network provider which serves as a data container, while a complex client side architecture performs the access control. The server only provides availability and client is responsible for data confidentiality and integrity. In addition to content data, the architecture is also able to protect the social graph information.

Decentralized Architectures: As described earlier, Persona [7] is the state-of-the-art decentralized architecture for social network privacy. EASiER is similar to Persona as both use ABE, but EASiER also provides an efficient mechanism for revocation, thereby avoiding the overhead of re-keying with group members as well as re-encryption of old data.

In the approach by Shakimov et al. [35], users store their data in a Virtual Independent Server (VIS) owned by themselves. These VISs form an overlay network per OSN. The authors consider different types of decentralized OSNs, depending on where the

VISs reside: a) cloud b) desktop , and c) hybrid based, and compare their privacy, costs and availability. Diaspora [21] is a decentralized OSN that users install on their own personal web servers. Backes et al. [36] present a core API for social networking, which can also constitute a plug-in for distributed OSNs. They assume that the server is trusted with the data while implementing access control. Both of these approaches avoid encryption.

PeerSon [37], LotusNet [38] and Safebook [39], three decentralized designs for social networking benefit from DHTs in their architecture. PeerSon and Safebook suggest access control through encryption, but they fall short in providing fine-grained policies comparing to ABE-based access control. Safebook is based on a peer-to-peer overlay network named *Matryoshka*. The end-to-end privacy in Matryoshka is provided by leveraging existing hop-by-hop trust. In all of these schemes overhead of key revocation affects performance.

8 CONCLUSION

We presented a scheme called PIRATTE for efficient revocation in *Ciphertext Policy Attribute-based Encryption*. We achieved this revocation scheme by introducing a semi-trusted proxy, leveraging ideas from a group communication scheme, and combining it with ABE. We also presented an architecture named EASiER for Online Social Networks (OSNs) that enforces access control through encryption using techniques in PIRATTE. Although we showed the use of PIRATTE in an OSN setting, it can be applied to any context where ABE is used for data protection with dynamic group membership. We implemented the scheme and compared it with Bethencourt et al.'s CP-ABE. Our results show that PIRATTE is scalable in terms of computation and communication for OSNs; accordingly, we have built a prototype application in the Facebook OSN to provide such encryption.

ACKNOWLEDGMENTS

The authors would like to thank Prateek Mittal.

REFERENCES

- [1] M. O'Connor, "Facebook Revealed Private Email Addresses Last Night," in *GAWKER*, 2010.
- [2] P. Wong, "Conversations About the Internet #5: Anonymous Facebook Employee," *The Rumpus*, 2010.
- [3] A. Sahai and B. Waters, "Fuzzy identity based encryption," in *Eurocrypt*, 2005.
- [4] J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext-Policy Attribute-Based Encryption," in *IEEE S & P*, 2007.
- [5] V. Goyal, O. Pandey, A. Sahai, and B. Water, "Attribute-based Encryption for Fine-grained Access Control of Encrypted Data," in *ACM CCS*, 2006.
- [6] S. Yu, C. Wang, K. Ren, and W. Lou, "Attribute Based Data Sharing with Attribute Revocation," in *ASIACCS*, 2010.
- [7] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin, "Persona: An Online Social Network with User-Defined Privacy," in *ACM SIGCOMM*, 2009.
- [8] S. Guha, K. Tang, and P. Francis, "NOYB: Privacy in Online Social Networks," in *WOSN*, 2008.
- [9] M. M. Lucas and N. Borisov, "FlyByNight: Mitigating the Privacy Risks of Social Networking," in *WPES*, 2008.
- [10] M. Naor and B. Pinkas, "Efficient Trace and Revoke Schemes," in *FC*, 2001.
- [11] A. Shamir, "How to Share a Secret," *CACM*, vol. 22, no. 11, 1979.
- [12] A. Miyaji, M. Nakabayashi, and S. TAKANO, "New Explicit Conditions of Elliptic Curve Traces for FR-Reduction," 2001.
- [13] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer, "Social Phishing," *CACM*, vol. 50, no. 10, 2007.
- [14] M. Cha, A. Mislove, B. Adams, and K. P. Gummadi, "Characterizing Social Cascades in Flickr," in *WOSN*, 2008.
- [15] R. Gross and A. Acquisti, "Information Revelation and Privacy in Online Social Networks (The Facebook case)," in *WPES*, 2005.
- [16] W. Luo, Q. Xie, and U. Hengartner, "FaceCloak: An Architecture for User Privacy on Social Networking Sites," in *PASSAT*, 2009.
- [17] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and Analysis of Online Social Networks," in *IMC*. ACM, 2007.
- [18] S. Jahid, P. Mittal, and N. Borisov, "EASiER: Encryption-based Access Control in Social Networks with Efficient Revocation," in *ASIACCS*, 2011.
- [19] S. Jahid, S. Nilizadeh, P. Mittal, N. Borisov, and A. Kapadia, "DECENT: A Decentralized Architecture for Enforcing Privacy in Online Social Networks," in *SESOC*, 2012.
- [20] "CP-ABE Toolkit," <http://acsc.cs.utexas.edu/cpabe/>.
- [21] D. Grippi, M. Salzberg, R. Sofaer, and I. Zhitomirskiy, "DIASPORA*," <https://joindiaspora.com/>.
- [22] N. Attrapadung and H. Imai, "Conjunctive Broadcast and Attribute-Based Encryption," in *Pairing*, 2009.
- [23] A. Boldyreva, V. Goyal, and V. Kumar, "Identity-based Encryption with Efficient Revocation," in *ACM CCS*, 2008.
- [24] R. Ostrovsky, A. Sahai, and B. Waters, "Attribute-based Encryption with Non-monotonic Access Structures," in *ACM CCS*, 2007.
- [25] A. Lewko, A. Sahai, and B. Waters, "Revocation Systems with Very Small Private Keys," in *IEEE S & P*, 2010.
- [26] J. Hur, "Improving Security and Efficiency in Attribute-Based Data Sharing," *IEEE TKDE*, 2011.
- [27] M. Blaze, G. Bleumer, and M. Strauss, "Divertible Protocols and Atomic Proxy Cryptography," in *EUROCRYPT*, 1998.
- [28] G. Ateniese, K. Fu, M. Green, and S. Hohenberger, "Improved Proxy Re-encryption Schemes with Applications to Secure Distributed Storage," in *NDSS*, 2005.
- [29] R. Canetti and S. Hohenberger, "Chosen-ciphertext Secure Proxy Re-encryption," in *ACM CCS*, 2007.
- [30] B. Libert and D. Vergnaud, "Unidirectional Chosen-ciphertext Secure Proxy Re-encryption," in *PKC*, 2008.
- [31] M. Green and G. Ateniese, "Identity-Based Proxy Re-encryption," in *ACNS*, 2007.
- [32] X. Liang, Z. Cao, H. Lin, and J. Shao, "Attribute Based Proxy Re-encryption with Delegating Capabilities," in *ASIACCS*, 2009.
- [33] K. Singh, S. Bhola, and W. Lee, "xbook: Redesigning privacy control in social networking platforms," in *USENIX Security*, 2009.
- [34] J. Anderson, C. Diaz, J. Bonneau, and F. Stajano, "Privacy-Enabling Social Networking Over Untrusted Networks," in *WOSN*, 2009.
- [35] A. Shakimov, A. Varshavsky, L. Cox, and R. Caceres, "Privacy, Cost, and Availability Tradeoffs in Decentralized OSNs," *WOSN*, 2009.
- [36] M. Backes, M. Maffei, and K. Pecina, "A security API for distributed social networks," in *NDSS*, 2011.
- [37] S. Buchegger, D. Schiöberg, L. H. Vu, and A. Datta, "PeerSoN: P2P Social Networking — Early Experiences and Insights," in *SNS*, 2009.
- [38] L. M. Aiello and G. Ruffo, "LotusNet: Tunable Privacy for Distributed Online Social Network Services," *Computer Communications*, 2012.
- [39] L. A. Cuttillo, R. Molva, and T. Strufe, "Safebook: Feasibility of Transitive Cooperation for Privacy on a Decentralized Social Network," in *WOWMOM*, 2009.